

**ADAPTATION PERFORMANCE IMPROVEMENTS FOR
ACTIVE CONTROL OF SOUND OR VIBRATION**

- [1] This application claims priority to U.S. Provisional Application Serial No. 60/271,470,
Filed February 27, 2001.

BACKGROUND OF THE INVENTION

- [2] Field of the Invention

This invention relates to optimal control of a system such as an active noise or vibration control system such as for a helicopter.

- [3] Background

An active control system consists of a number of sensors which measure the ambient variables of interest (e.g. sound or vibration), a number of actuators capable of generating an effect on these variables (e.g. by producing sound or vibration), and a computer which processes the information received from the sensors and sends commands to the actuators so as to reduce the amplitude of the sensor signals. The control algorithm is the scheme by which the decisions are made as to what commands to the actuators are appropriate.

- [4] For tonal control problems, the computation can be performed at an update rate lower than the sensor sampling rate as described in copending application entitled "Computationally Efficient Means for Active Control of Tonal Sound or Vibration". This approach involves demodulating the sensor signals so that the desired information is near DC (zero frequency), performing the control computation, and remodulating the control commands to obtain the desired output to the actuators. The control computations are therefore performed on the sine and

cosine components at the frequency of interest for each sensor signal. These can be represented as a complex variable where the real part is equal to the cosine term, and the imaginary part is equal to the sine term.

[5] The number of sensors is given by n_s and the number of actuators is n_a . The complex harmonic estimator variables that are calculated from the measurements of noise or vibration level can be assembled into a vector of length n_s denoted z_k at each sample time k . The control commands generated by the control algorithm can likewise be assembled into a vector of length n_a denoted u_k . The commands sent to the actuators are generated by multiplying the real and imaginary parts of this vector by the cosine and sine of the desired frequency.

[6] In the narrow bandwidth required for control about each tone, the transfer function between actuators and sensors is roughly constant, and thus, the system can be modeled as a single quasi-steady complex transfer function matrix, denoted T . This matrix of dimension n_s by n_a describes the relationship between a change in control command and the resulting change in the harmonic estimate of the sensor measurements, that is, $\Delta z_k = T \Delta u_k$. For notational simplicity, define $y_k = \Delta z_k$, and $v_k = \Delta u_k$. The complex values of the elements of T are determined by the physical characteristics of the system (including actuator dynamics, the structure and/or acoustic cavity, and anti-aliasing and reconstruction filters) so that T_{ij} is the response at the reference frequency of sensor i due to a unit command at the reference frequency on actuator j . Many algorithms may be used for making control decisions based on this model.

[7] The control law is derived to minimize a quadratic performance index

$$J = z^T W_z z + u^T W_u u + v^T W_{\delta u} v$$

where W_z , W_u and $W_{\delta u}$ are diagonal weighting matrices on the sensor, control inputs, and rate of

change of control inputs respectively. A larger control weighting on an actuator will result in a control solution with smaller amplitude for that actuator.

[8] Solving for the control which minimizes J yields:

$$u_{k+1} = u_k - Y_k(W_u u_k + T_k^T W_z z_k)$$

where

$$Y_k = (T_k^T W_z T_k + W_u + W_{\delta u})^{-1}$$

[9] The matrix Y determines the rate of convergence of different directions in the control space, but does not affect the steady state solution. In the following equation, the step size multiplier $\beta < 1$ provides control over the convergence rate of the algorithm. A value of approximately $\beta = 0.1$ may be used, for example.

$$u_{k+1} = u_k - \beta Y_k(W_u u_k + T_k^H W_z z_k)$$

[10] The performance of this control algorithm is strongly dependent on the accuracy of the estimate of the T matrix. When the values of the T matrix used in the controller do not accurately reflect the properties of the controlled system, controller performance can be greatly degraded, to the point in some cases of instability. An initial estimate for T can be obtained prior to starting the controller by applying commands to each actuator and looking at the response on each sensor. However, in many applications, the T matrix changes during operation. For example, in a helicopter, as the rotor rpm varies, the frequency of interest changes, and therefore the T matrix changes. For the gear-mesh frequencies, variations of 1 or 2% in the disturbance frequency can result in shifts through several structural or acoustic modes, yielding drastic phase and magnitude changes in the T matrix, and instability with any fixed-gain controller (i.e. if $T_{k+1} = T_k$ for all k). Other sources of variation in T include fuel burn-off, passenger movement,

altitude and temperature induced changes in the speed of sound, etc.

- [11] There are several possible methods for performing on-line identification of the T matrix, including Kalman filtering, an LMS approach, and normalized LMS. For an estimated T matrix, T^e , an error vector can be formed as

$$E = y - T^e v$$

- [12] The estimated T matrix is updated according to

$$T_{k+1}^e = T_k^e + E K^T$$

- [13] The different estimation schemes differ in how the gain matrix K is selected. The Kalman filter gain K is based on the covariance of the error between T and the estimate T^e , given by the matrix P where

$$M = P_k + Q$$

$$K = Mv / (R + v^T Mv)$$

$$P_{k+1} = M - Kv^T M$$

and the matrix Q is a diagonal matrix with the same dimension as the number of actuators, and typically with all diagonal elements equal. The scalar R can be set equal to one with no loss in generality provided that the matrices Q and R are constant in time. The normalized LMS approach is very similar, with the gain matrix K given by

$$K = Qv / (1 + v^T Qv)$$

- [14] The algorithm can be used with the Kalman filter approach, or using the normalized LMS approach which is computationally simpler and may provide similar or better performance. The current invention is described in terms of this equation, however, the specific form of the adaptation algorithm is not crucial to the invention.

[15] Any of these adaptation schemes will obtain excellent estimation of the T matrix when there is little noise in the measurements. As noise levels increase, however, there are difficulties as the filter can not distinguish between the effects of noise and the effects of actual changes in T. As a result, the adaptation parameters will tend to drift. Decreasing adaptation gain will decrease the drift but not prevent it, and will degrade the adaptation performance.

[16] The algorithm as described above is self-adaptive in the sense that the plant (i.e., system) model used in the control update calculation is actively updated during closed-loop operation based on changes in the sensor signals resulting from the application of the changes in the actuator commands determined by the control update calculation. However, during steady-state conditions, when changes in the control commands are responding only to "noise" in the estimate of the disturbance being canceled, there is a loss of identifiability of the plant. This loss of identifiability is a result of coupling the control update calculation and the adaptation update calculation together. The adaptation process is estimating the system by observing how changes in actuator commands cause changes in the measured system response; however, due to the control process, the change in control Δu is not independent from the change in measurement Δz . This coupling results in an instability or drift observed in steady state that can be severe if signal-to-noise is poor.

[17] In addition to causing drift behavior, as described above, noise also degrades the quality of the adaptation estimate through two related mechanisms. First, because the signal to noise ratio is reduced the adaptation estimate will include a random noise component. The second effect is more subtle. The adaptation is intended to estimate the transfer function from actuators to sensors at the disturbance frequency. It is reasonable to assume that the physical transfer

function does not vary substantially over the bandwidth of the control. Control transients are close to the disturbance frequency, and therefore the transfer functions due to these input signals are close to the desired transfer function. However, the adaptation will also respond to noise on y_k and v_k that is at higher frequencies (in the demodulated system, i.e. further away from the desired frequency). Changes in the physical system at these frequencies are unknown, however, the most significant change is due to the harmonic estimation filter on y_k , the effect of which is known. The estimation of the harmonic components of the sensor signals at the desired frequency requires low-pass filtering to avoid aliasing in downsampling the estimates. While there is no reduction in the amplitude of the sensor information y_k in the immediate vicinity of the desired frequency, y_k is attenuated at frequencies further from the desired frequency, while v_k is not. The algorithm described above will therefore under-predict the desired transfer function because it will average over frequency, and include part of the reduction in amplitude of the harmonic estimator filter. There will be no bias in phase since the noise for positive frequencies should be comparable to the noise for negative frequencies, and the phase effects will therefore cancel.

SUMMARY OF THE INVENTION

- [18] The present invention improves the quality of the adaptation estimates by filtering the input signals to the adaptation, selectively implementing a “dead-zone” during which adaptation does not occur and by selectively adding a dither signal to the control commands. In the present invention, the dead-zone is based on the magnitude of the changes in the control commands, which are responding to the changes in the sensor signals. The dead-zone can be applied to all actuators simultaneously, or can be applied to the adaptation of each actuator channel

independently. The appropriate threshold for the dead-zone can be updated on-line.

[19] To maintain identifiability, a “dither signal” is added to the control commands to “ping” the system to increase the amount of information available to the adaptive algorithm. The dither signal is preferably implemented on only one actuator channel at a time. Further, it is preferred that triangular dither is implemented such that the control amplitude of the channel being dithered is ramped down and then ramped back up to its initial value in a predetermined number of steps. Also, the dither amplitude for each channel is preferably set to be proportional to the current control amplitude.

[20] This invention will also improve the quality of the adaptation estimates by filtering the input signals to the adaptation circuit, $y_k = \Delta z_k$, and $v_k = \Delta u_k$. Among other things, the filter matches the filtering applied by the harmonic estimator.

BRIEF DESCRIPTION OF THE FIGURES

[21] FIG. 1 shows a block diagram of the noise or vibration control system of the present invention.

[22] FIG. 2 shows a vehicle in which the present invention may be used.

DETAILED DESCRIPTION

[23] Control systems consist of a number of sensors which measure ambient vibration (or sound), actuators capable of generating vibration at the sensor locations, and a computer which processes information received from the sensors and sends commands to the actuators which generate a vibration field to cancel ambient vibration (generated, for example by a disturbing

force at the helicopter rotor). The control algorithm is the scheme by which the decisions are made as to what the appropriate commands to the actuators are.

[24] FIG. 1 shows a block diagram 10 of an active control system. The system comprises a structure 102, the response of which is to be controlled, sensors 128, filter 112, control unit 107 and actuators (which could be force generators) 104. A disturbance source 103 produces undesired response of the structure 102. In a helicopter, for example, the undesired disturbances are typically due to vibratory aerodynamic loading of rotor blades, gear clash, or other source of vibrational noise. A plurality of sensors 128(a)...(n) (where n is any suitable number) measure the ambient variables of interest (e.g. sound or vibration). The sensors (generally 128) are typically microphones or accelerometers, or virtually any suitable sensors. Sensors 128 generate an electrical signal that corresponds to sensed sound or vibration. The electrical signals are transmitted to filter 112 via an associated interconnector 144(a)...(n) (generally 144). Interconnector 144 is typically wires or wireless transmission means, as known to those skilled in the art.

[25] Filter 112 receives the sensed vibration signals from sensors 128 and performs filtering on the signals, eliminating information that is not relevant to vibration or sound control. The output from the filter 112 is transmitted to control unit 107, which includes adaptation circuit 108 and controller 106, via interconnector 142. In the present invention, a filter 109 is placed before adaptation circuit 108, as will be described below. The controller 106 generates control signals that control force generators 104(a)...(n).

[26] A plurality of actuators 104(a)...(n) (where n is any suitable number) are used to generate a force capable of affecting the sensed variables (e.g. by producing sound or vibration). Force

generators 104(a)...(n) (generally 104) are typically speakers, shakers, or virtually any suitable actuators. Actuators 104 receive commands from the controller 106 via interconnector 134 and output a force, as shown by lines 132(a)...(n) to compensate for the sensed vibration or sound produced by vibration or sound source 103.

[27] The control unit 107 is typically a processing module, such as a microprocessor. Control unit 107 stores control algorithms in memory 105, or other suitable memory location. Memory 105 is, for example, RAM, ROM, DVD, CD, a hard drive, or other electronic, optical, magnetic, or any other computer readable medium onto which is stored the control algorithms described herein. The control algorithms are the scheme by which the decisions are made as to what commands to the actuators 104 are appropriate, including those conceptually performed by the controller 107 and adaptation circuit 108. Generally, the mathematical operations described in the Background, as modified as described below, are stored in memory 105 and performed by the control unit 107. One of ordinary skill in the art would be able to suitably program the control unit 107 to perform the algorithms described herein. The output of the adaptation circuit 108 may be filtered before being sent to the controller 107.

[28] The present invention improves the quality of the adaptation estimates by adaptation circuit 108 by filtering the input signals to the adaptation circuit, selectively implementing a “dead-zone” during which adaptation does not occur and by selectively adding a dither signal to the control commands.

[29] The filtering, dead zone and dither can be used independently; however, dead zone is beneficial to overcome the drift behavior and dither is beneficial to maintain good performance while using dead zone.

[30] Dead zone

One way to avoid the aforementioned loss of identifiability during steady-state conditions is to actively interrogate the system to determine if the operating conditions are fairly steady and if so turn adaptation off, and only turn it back on again when the system is no longer in a steady-state condition. This approach is usually referred to as establishing a “dead-zone” in which adaptation does not occur. Usually, this dead-zone is established based on examining the magnitude of changes in the sensor signals from one iteration to the next, and if they are smaller than some predefined level, turning adaptation off, and if above that level, turning adaptation on.

In the present invention, the dead-zone is based instead on the magnitude of the changes in the control commands, which are responding to the changes in the sensor signals. The variable Δu provides a good estimate of whether there are any significant changes that require adaptation to be turned on. Any change in the system T matrix will result in changes in Δu . The use of $\|\Delta u\|$ as the dead-zone switch is motivated by the fact that there is some minimum value of $\|\Delta u\|$ required for each channel to achieve a good signal-to-noise ratio.

[31] The dead-zone can be applied to all actuators simultaneously, or can be applied to the adaptation of each actuator channel independently (i.e. to each column of the T matrix independently). With the former, then every element of the diagonal matrix Q is equal and set according to $q = q_0$ if $\|\Delta u\| > \Delta_z$ and $q = 0$ otherwise, where Δ_z is the dead-zone threshold. If the dead-zone is instead applied to each channel, then each element of the diagonal matrix Q is either zero if the corresponding $|\Delta u_i| < \Delta_z$, or set equal to q_0 otherwise. Setting $Q(i)=0$ turns adaptation off for channel i, and setting it to some positive number greater than zero turns adaptation on. Performing the dead-zone comparison on an actuator channel-by-channel basis is preferable.

[32] Rather than switching adaptation only between on and off, the adaptation gain can be any function of $|\Delta u|$ that is small for small $|\Delta u|$, some constant for large $|\Delta u|$, and goes through some transition in between. For example, setting $q = \min(q_0, q_v |\Delta u|^N)$ has been tested, where N is an integer typically between 2 and 6. Again, this can be set independently for each actuator channel, or the absolute value can be replaced by a norm of the overall vector and the adaptation gain set equal for all channels.

[33] A further refinement of this approach is to actively adjust or recalculate the dead-zone threshold for each actuator channel based on the current noise floor and actuator effectiveness. The dead-zone threshold for each channel is obtained by slowly filtering the dead-zone parameter $|\Delta u_i|$ and multiplying the result by a scale factor typically between 2 and 3. The filter time constant can be 100 times larger than the adaptation time constant.

[34] Dither

The second aspect of this invention is to ensure that the adaptation algorithm has sufficient information to accurately identify the system. To maintain identifiability, a “dither signal” is added to the control commands to, in essence, “ping” the system to increase the amount of information available to the adaptive algorithm. Although it is known to add a dither signal to this type of system, there are several unique aspects of the dither signal implementation in the present invention.

[35] First, the dither signal is implemented on only one actuator channel at a time. Because only a single channel is dithered, then a larger amplitude can be used without degrading overall performance significantly, and hence a larger signal-to-noise ratio is obtained. All of the information coherent with the injected dither signal can be associated with the dithered actuator

104 channel.

[36] The channel being dithered is removed from the control update calculation with a high control weighting, while being the only channel included in the adaptation update calculation. This modification decouples the control and adaptation behavior so that the adaptation recovers its ideal open-loop behavior, and avoids any possibility of drift. Furthermore, any sensor information resulting from the dither will not corrupt the system transfer function estimates for any of the other actuators. The system looks at the response of the sensors 128 taking into account the fact that the actuators 104, other than the dithered actuator, have compensated for the dither signal added to the dithered actuator.

[37] Triangular dither is implemented such that the control amplitude of the channel being dithered is ramped down and then ramped back up to its initial value in a predetermined number of steps. The motivation for ramping the control amplitude down first and then back up to its original value is that in the case where the dithered channel is up against the saturation limit (maximum allowable actuator command) before the application of dither, then the dither will bring the actuator 104 down off this limit. Alternatively, if the dither were implemented first as a ramp up in amplitude followed by a ramp down, then in the case of saturation, control limiting (i.e., signal clipping) would prevent the dither from being implemented since the actuator 104 would be up against the saturation limit throughout the entire dither cycle. Other types of dither signal are possible; however, the ramp maximizes the variable Δu which is used by the adaptation algorithm.

[38] The dither amplitude for each channel is set to be proportional to the current control amplitude. This is desirable because not all actuators have equal authority. An actuator 104 with

higher authority requires lower control amplitude to achieve the same output or signal to noise ratio. Using this improvement will therefore result in a lower dither amplitude for such an actuator, which, due to the higher authority, will result in similar signal-to-noise for all actuators.

Without this modification, a fixed dither amplitude for every actuator 104 would require that a large amplitude be used for sufficient signal-to-noise on the lower authority actuators, which would then yield significant performance degradation of the control system while dithering the higher authority actuators. An additional benefit of a proportional dither signal is to provide more information in highly uncertain directions. If some information in the estimate of the T matrix has more error than other information, this will tend to result in larger than necessary amplitudes on the control signals corresponding to the bad information. As a result, the dither amplitude will increase for these control signals, providing more information to correct the poorer T matrix estimates. If the estimate of the T matrix is known to be poor, for example when the active control system is first activated, then a more rapid adaptation convergence can be obtained by temporarily increasing the dither amplitude and/or the adaptation gain.

[39] This invention will also improve the quality of the adaptation estimates by filtering the input signals to the adaptation circuit 108, $y_k = \Delta z_k$, and $v_k = \Delta u_k$ with filter 109. The presence of background noise not only results in parameter drift, but results in both a bias and a random error in the elements of the T matrix estimate. Improving the quality of the estimate improves the achievable performance of the control algorithm. Alternately, the same quality estimate can be achieved, and the invention can be used to allow lower amplitude dither signals to obtain similar adaptation performance, and lower degradation in overall performance due to the dither.

[40] Introducing a matching first-order low-pass filter on the signal v_k corrects for the bias

effect caused by the adaptation responding to noise which has been filtered through the harmonic estimator. This matching filter must have the same corner frequency as the harmonic estimator filter. The result is improved signal to noise due to lower emphasis on high frequency noise, and eliminating bias in the estimate.

[41] The adaptation algorithm described in the background uses the difference in z and u between samples, $\Delta z = z_k - z_{k-1}$ and $\Delta u = u_k - u_{k-1}$, in order to avoid variations in the ambient near zero frequency from distorting the transfer function estimate. This differencing amplifies higher frequency content in u and z . The combination of this with the harmonic estimator filter on z and its matching filter on u results in all information at frequencies above the harmonic estimator corner frequency being treated with equal importance, and information below this frequency being treated with lower importance. Since the information at higher frequencies is predominantly noise rather than useful information, including additional low-pass filtering on both the signals v_k and y_k will improve signal to noise ratio, and improve the quality of the T -matrix estimate (or conversely, allow the same quality of estimate with a lower amplitude dither signal, and hence improve overall algorithm noise reduction performance.) This low-pass filtering can be implemented either as a separate step, or by changing the difference definition of v_k and y_k to a high-pass filter. The result of the combination of filters is to put the most emphasis on the information at the frequencies around the filter corner frequencies, and less emphasis on the information at lower and higher frequencies.

[42] A further modification of the above filtering could improve adaptation performance by using a separate set of filters on the information generated by control transients, and the information generated by the dither signal. The filter characteristics can then be optimized for

the known frequency domain characteristics of each of the two types of control inputs u_k , in order to maximize the signal to noise ratio for each.

[43] A final modification to the above concept could be useful in systems with small numbers of actuators. The dither algorithm contained in this invention uses time separation to distinguish between the information on one actuator 104 and the information on another, and this is appropriate if there are many actuators. If there are very few actuators, then frequency separation could be used instead. That is, one could put a continuous, small amplitude dither signal on each actuator command, with a different frequency for each actuator. Second-order resonant filters on both v_k and y_k could then be used to extract the distinct information corresponding to each frequency.

[44] FIG. 2 shows a perspective view 20 of a vehicle 118 in which the present invention can be used. Vehicle 118, which is typically a helicopter, has rotor blades 119 (a)...(d). Gearbox housing 110 is mounted at an upper portion of vehicle 118. Gearbox mounting feet 140 (a)...(c) (generally 140) provide a mechanism for affixing gearbox housing 110 to vehicle airframe 142. Sensors 128(a) through (d) (generally 128) are used to sense acoustic vibration produced by the vehicle, which can be from the rotorblades 119 or the gearbox housing 110. Although only four sensors are shown, there are typically any suitable number of sensors necessary to provide sufficient feedback to the controller (not shown). The sensors 128 may be mounted in the vehicle cabin, on the gearbox mounting feet 140, or to the airframe 142, or to another location on the vehicle 118 that enables vehicle vibrations or acoustic noise to be sensed. Sensors 128 are typically microphones, accelerometers or other sensing devices that are capable of sensing vibration produced by gear clash from the gearbox 110 and generating a signal as a function of

the sensed vibration. These sensors generate electrical signals (voltages) that are proportional to the local noise or vibration.

[45] In accordance with the provisions of the patent statutes and jurisprudence, exemplary configurations described above are considered to represent a preferred embodiment of the invention. However, it should be noted that the invention can be practiced otherwise than as specifically illustrated and described without departing from its spirit or scope. Alphanumeric identifiers for steps in the method claims are for ease of reference by dependent claims, and do not indicate a required sequence, unless otherwise indicated.